

Pension taxation, household debt and the real economy

Af
Henrik Yde Andersen*

This version: 23 October 2018

Keywords: Crowding out, retirement savings, household debt
JEL codes: D14, D15, E21, H31

Abstract

This paper reviews findings from the literature about the interplay between pension savings and debt accumulation in households. The first part explains why debt levels in the household sector play a role for macroeconomic developments. The second part presents how pension taxation affects the pace at which households accumulate debt. One recent contribution in this field is Andersen (2018) which provides an empirical analysis on crowding out in retirement accounts. The results show that savers tend to redirect savings from pension accounts towards alternative savings accounts once tax incentives for saving in retirement accounts are reduced. Specifically, 1/3 of savings that was usually placed in tax-favored pension accounts was now used for debt repayment, while 2/3 was placed in other savings accounts. The findings imply that pension taxation may be an important factor when explaining household debt accumulation.

* Danmarks Nationalbank (hya@nationalbanken.dk). Based on a presentation prepared for Danish Economic Society's biennial meeting at Hotel Koldingfjord 2018. Thanks to meeting participants as well as colleagues at Danmarks Nationalbank for comments. Viewpoints and conclusions stated are the responsibility of the author, and do not necessarily reflect the views of Danmarks Nationalbank. Remaining errors are the author's responsibility alone.

1. Introduction

Over the past decade, increased attention have been devoted to whether household debt levels affect real economic outcomes. Highly indebted households are committed to pay a large share of their income to service the debt. Compared to their less leveraged peers these households are vulnerable to unanticipated shocks to interest rates and may need to cut consumption if interest rates increase, and ultimately, rate hikes could force them to default on their debt.

Further, a sudden drop in asset prices could undermine the collateral posted for their liabilities, for example if the home value falls below the value of outstanding debt. To avoid this, borrowers could choose to speed up debt repayments. If they have no other assets to draw down this could also force homeowners to cut spending. These mechanisms imply that consumption is likely to be increasingly volatile with high debt levels in the household sector, and may potentially prolong the recovery period after an economic recession.

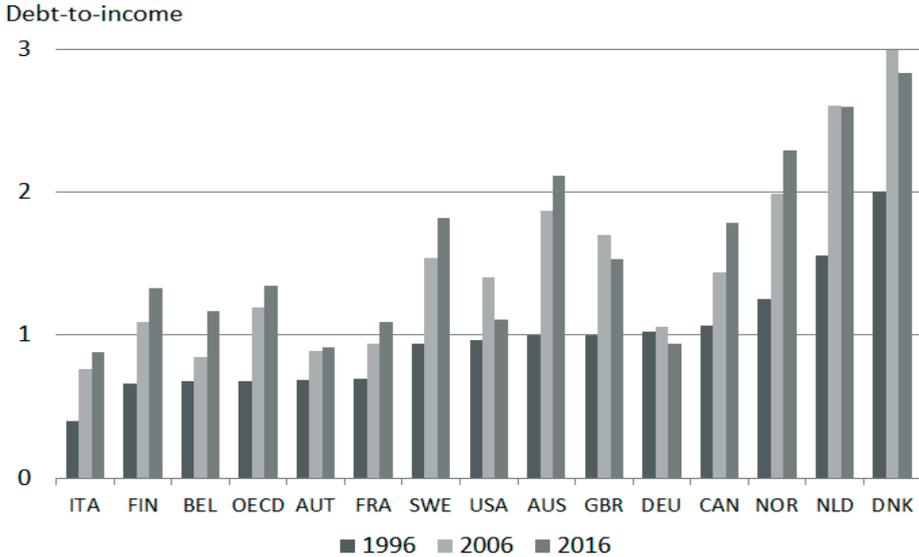
High debt in households could, however, be a natural consequence of institutional settings. Cross-country evidence shows that household saving rates, including savings in retirement accounts, are correlated with household debt levels. This implies that liabilities are counterbalanced, so to speak, by assets such that the net debt remains low. In this case, borrowers will be less exposed to changes in interest rates and the effect of high debt levels on consumption should be less important.

Household debt ratios – that is, total debt divided by disposable income – has risen in most developed countries for the past decades, indicating that debt increases at a higher pace than income. Figure 1 shows household debt-to-income ratios for a range of selected OECD countries in 1996, 2006 and 2016. There are a few points to mention about these developments. First, there is substantial variation across countries in household debt ratios. Second, debt ratios have increased for almost all countries over time. Only Germany has fairly unchanged household leverage. Third, USA and Great Britain – and to some extent Denmark – show signs of post-crisis deleveraging. Cross-country variation in institutional settings is likely to explain the observed variation in household leverage. We therefore turn to the case of Denmark, which consistently has the highest household leverage ratio.

Based on Danish administrative data, Andersen (2018) argues that beneficial tax treatment of pension savings affects household debt accumulation. This could be the case if the net return from investing in a pension plan is higher than the payoff from repaying outstanding debt. To test this hypothesis Andersen (2018) designs a quasi-experiment to quantify the change in debt repayments that was caused by a sudden reduction in after-tax returns on pension contributions. Such interplay between pension taxation and changes in household debt could explain

some of the empirical patterns observed in advanced economies over the past decades, in particular the balance sheet expansion where both financial assets and liabilities increase with little or no effect on the net debt position. Another – and potentially more important – explanation of the balance sheet expansion is the introduction of mandatory pension contribution. This element is, however, not addressed in this paper.

Figure 1: Household debt-to-income in selected OECD countries



Note: The graph shows household debt relative to net disposable income in 1996, 2006 and 2016 for a range of selected OECD countries, including the OECD average.

Source: OECD

Section 2 reviews the literature, first on the link between household debt levels and macroeconomic outcomes and secondly on the effects of tax incentives on savings behavior. Section 3 discusses the interplay between pensions and household debt accumulation. Finally, section 4 concludes.

2. Literature

This section reviews papers that provided important insights and relevant discussions about the role of household debt on macroeconomic developments and whether tax incentives for saving in retirement accounts are useful to stimulate savings rates.

2.1. Household debt and the real economy

Eggertsson & Krugman (2012) were among the first to formalize the interaction between debt levels and spending in a structural economic model. By developing a DSGE-model with two types of consumers, impatient and patient ones, the authors were able to describe the transmission mechanism between household leverage and real economic outcomes. The impatient agents preferred to consume now and thus borrow for current spending at the cost of lower spending opportunities in the future. The patient ones earned more than they wanted to spend today, enabling them to lend money to the impatient consumers. One could think of the two agent types as younger (impatient) and older (patient) consumers. The model imposes a borrowing limit on the impatient households. This limit dictates that the impatient households can only borrow up to a certain value of their assets, i.e. a loan-to-value constraint. The now credit constrained households must reduce spending and debt in order to comply with the borrowing limit, consistent with the idea that highly indebted consumers start deleveraging at an increased pace. The key contribution of the paper is the demonstration of the transmission of high debt levels to spending cuts by households. In other words, the low economic activity in the post-crisis years could potentially be linked to the pre-crisis debt levels in households, even when other general equilibrium components are accounted for. This lower economic activity arises because the impatient households are forced to pay down debt in response to declining asset prices, limiting the impatient households' ability to post collateral for borrowing. This in turn depresses aggregate demand, which feeds back to lower asset prices and a vicious loop is set in motion. The loop is stronger if the monetary policy rate hits the effective lower bound, limiting the central bank's possibility of stimulating the economy. The findings are important because they describe the interplay between the financial sector and macroeconomic developments in a theory consistent framework. The question is, however, whether the model predictions can be identified empirically.

In a cross-country study that includes 14 developed countries, Jorda *et al.* (2013) compares more than 200 recessions and finds that economic downturns caused by financial crises typically last longer than other types of recessions. Also, Cecchetti *et al.* (2011) finds that moderate levels of household debt enhances growth, while debt above some certain threshold is harmful to economic growth. In other words, high levels of debt create volatility according to the authors. Their study investigates the correlation between economic growth and debt levels in households, firms and governments in a range of OECD countries. Common for these two studies is that they examine financial conditions that potentially affect macroeconomic developments using macro-level data, e.g. national accounts. In order to gain a deeper understanding of the mechanisms that connect household debt levels to the decline in private consumption, more disaggregated evidence of the phenomenon is needed.

Dynan (2012) uses information at the household-level from a US survey to show that families with high debt levels prior to the Great Recession reduced spending significantly in the years following the recession relative to families without *ex ante* large debt positions.¹ This study is important as it provides direct evidence of the existence of a link between household debt and the marginal propensity to consume. The advantage of using microdata is that the author is able to control for a range of family characteristics. This implies that differences in, e.g., income, wealth and age do not explain the observed spending cut. There is, however, one crucial caveat to the study. The author is unable to reject the possibility of reversed causality, implying that the decline in private consumption could have led to a decrease in the demand for credit. This endogeneity problem is a key challenge to these types of studies. In order to claim causality, namely that highly leveraged households were more exposed to the recession than their less leveraged peers, researchers need variation in household debt levels which is not systematically connected to the consumption pattern of the very same households.

Mian *et al.* (2013) attempt to deal with the innate endogeneity problem by exploiting variation in house prices across US states during the financial crisis. While some states saw a significant drop in house prices other states were less affected. The ones that experienced sharp price declines were also the ones that had massive house price increases in years immediately before the recession. The authors suggested that the states with high house price volatility across the business cycle could be characterized by limited supply of housing. The supply side of the housing market was more or less fixed in some states because the geographic conditions were unfavorable to building additional housing, e.g. rocks, sea or other poor conditions in the underground. In other states with ample space for new housing, there were no basis for significant fluctuations in house prices. Based on the argument that housing supply affects house prices directly but does not correlate with consumption and borrowing behavior, the study was able to identify a causal link between debt and consumption. Using information about how much debt the families in the survey had *ex ante* the authors could show that highly leveraged families reduced consumption significantly more than low leveraged households over the course of the recession. More recently, this pattern is found to prevail across a range of countries (Mian *et al.*, 2017).

1. The increase in household debt in decades prior to the outbreak of the 2007-2008 financial crisis is by some called the *democratization of credit* (Dynan, 2009). This term covers an expansion of financial opportunities across all income groups and not only by the wealthiest. Households might have borrowed more than they were able to service in the long run making them vulnerable to economic recessions. The author argues that the build-up of household debt prior to the crisis could be a destabilizing factor, resulting in more volatile economic developments.

A study by Andersen *et al.* (2016) suggests that the interplay between high debt levels and subsequent spending cuts reflects a mechanical effect rather than a causal one. Their argument is that households are likely to always cut spending in the years after taking up debt. Imagine for instance a household who buys a new car in one year and repays the car loan over the following 7 years. Both debt and spending increases immediately in the year of the purchase. For the subsequent 7 years spending will be substantially lower because this exact household does not buy a new car each year. Households who financed such purchase with debt immediately before the Great Recession are likely to be marked as highly leveraged households, while those who did the purchase a few years earlier might be marked as less leveraged households. The authors argue that it is not necessarily the recession that caused the spending cut by highly leveraged households but rather the timing of durable spending.

It is, however, difficult to rule out the evidence in favor of a causal link between household debt levels and spending cuts. Particularly as such link may explain why the 2007–2008 economic crisis lasted longer than earlier recessions. In order to improve the resilience against such economic downturn, economists strive to understand the observed build-up of unsustainably high debt levels in households. The following section argues how tax incentives for saving in pension accounts potentially affect household debt levels.

2.2. Savings and tax incentives

Pension systems have been part of the financial infrastructure in developed countries for many decades. The basic idea about these systems is to spur savings, to finance investments, sustain growth, and help households save adequately for retirement (Engen *et al.*, 1996). Economists have studied the effectiveness of pension accounts in terms of stimulating individual savings rates for just as long and their findings vary hugely. Some studies find that the introduction of tax-favored pension accounts causes savers to increase overall savings significantly such that 1 dollar saved in a retirement account increases savings by up to an equal amount (See e.g. Venti & Wise, 1990; Skinner & Feenberg, 1990; Poterba *et al.*, 1995, 1996; Hubbard & Skinner, 1996). In this case, the retirement policy achieves the desired effect as the taxpayers respond by saving more than they would otherwise have done in absence of the tax benefit. Other papers find that savings in retirement accounts are fully crowded out by reducing savings in non-retirement accounts (See e.g. Gale & Scholz, 1994; Engen *et al.*, 1996; Gale, 1998; Attanasio & DeLeire, 2002; Attanasio & Rohwedder, 2003; Benjamin, 2003; Engelhardt & Kumar, 2007; Chetty *et al.*, 2014; Andersen, 2018). In this case, 1 dollar saved in a tax-favored pension account is followed by up to 1 dollar reduction in other savings accounts, e.g., savings in stocks, bonds or deposits. Here, the tax benefit would be an expensive and ineffective tool to stimulate savings rates. Bernheim (2002) provides

a detailed and thorough review on the discrepancies of these studies but a few of them are relevant to mention here.

The fundamental framework used to assess the effectiveness of tax-favored retirement accounts is the life-cycle model, stating that consumers smooth marginal utility of consumption across their lifetime. Specifically, agents save while in working age and dis-save while retired. The theory posits that consumers at each point in time plan future consumption taking into account available resources and expected future income. This implies that their consumption paths are revised each time new information arrives, e.g., the introduction of tax-favored pension accounts. Hubbard (1984) is among the first to model such economic behavior in order to estimate if saving in retirement accounts actually raised overall savings.² He uses information from a cross-sectional US survey to show that contributions for tax-favored retirement accounts do increase individuals savings. In fact, the effect is found to increase with marginal tax rates, implying that such tax incentives have a larger effect on those who benefit the most from them—those with highest tax rates.

The discussion over the effect of tax incentives for saving in retirement accounts on savings rates continued during the 1980s and 1990s. Despite that many studies made important contributions to the literature, I will not mention all of them here. Rather, Engen *et al.* (1996) provide a review of why some of these studies might overstate the effectiveness of tax-favored pension accounts and highlight two challenges that the literature had failed to recognize. The first challenge is the total absence of the fact that individuals have different tastes for saving, implying that some people simply like to save, while others do not. This implies that previous studies may have overstated the effect of tax-favored pension accounts on total savings simply because the group of contributors is likely to have a higher preference for saving than the group of non-contributors. This illustrates one of the core caveats in the literature; it is challenging to find a reasonable counterfactual when assessing the effects of tax incentives on savings.

Ideally, we would like to compare two randomly formed groups of individuals such that they compare in characteristics and tastes for saving. Then we allow the one group to save in retirement accounts in order to obtain a tax benefit and see whether their total savings increase over time relative to the control group—the group of savers who had no access to tax-favored pension accounts. Such experimental setup would seem impossible and not least unethical to carry out. Therefore, economists turn to so called quasi-experimental research designs—or

2. Despite poor data availability in the 1980s the hump-shaped pattern of asset accumulation over the life-profile was heavily discussed. See e.g. King & Dicks-Mireaux (1982). Today, a range of competing hypotheses aim to explain the incompatible data pattern to the life-cycle theory, e.g., precautionary savings and bequest motives (Dynan *et al.*, 2002) or behavioral biases (Laibson, 1997).

natural experiments—that resemble the ideal design closely. A good example of this is Attanasio & DeLeire (2002), who develop a new strategy to account for the unobserved heterogeneity of savers. Rather than comparing behavior of contributors to non-contributors, they compare savings behavior of new contributors to continuing contributors. Specifically, the only difference between these groups is the timing of when they take up contributing to a tax-favored pension account. Given that both groups contribute it is reasonable to assume that they have an equally large taste for saving. Their findings suggest that the introduction of retirement accounts did not increase savings but led savers to reshuffle their financial assets, i.e. shifted savings from non-retirement accounts to tax-favored pension accounts. Effectively, tax incentives for saving in pension schemes provide the contributors with a tax break for doing what they would have done anyway. Note, however, that the treated savers—the new contributors—would always be starting their pension contributions at a later point in time than the non-treated savers—the continuing contributors. Therefore, they would not face the same conditions under which they decide to contribute for retirement schemes. In other words, there could be important underlying factors, e.g., income expectations, financial market developments, interest rates or institutional settings that would always be different between the treatment and control group, which their paper is unable to take account for. However, their findings highlight the importance of tastes for saving as stated by Engen *et al.* (1996).

Chetty *et al.* (2014) and Andersen (2018) attempt to take the empirical tests one step closer to the ideal experiment. The reason for this is twofold. First, both studies use administrative registers that are longitudinal in nature and hold objective information about personal characteristics, including income and wealth at the individual level. Second, they exploit plausibly exogenous variation in the tax subsidy obtained by contributing to pension accounts, which was caused by sudden tax policy changes. In combination, these are significant innovations compared to previous studies. The panel dimension of the data allows the authors to follow the behavior of the same individuals across time. The simple idea is to predict *ex ante* which savers are likely to be affected by the new tax policy changes and then follow their savings behavior across the implementation of the tax policy change. Chetty *et al.* (2014) exploit a Danish 1999 policy change that reduced the value of tax subsidies of pension contributions for savers above the top tax threshold. Using the discontinuity exactly around the top tax cutoff they are able to measure changes in savings rates and compare the affected savers (top-tax earners) with the unaffected ones (those just below the top-tax bracket). Andersen (2018) designs an empirical setup using a Danish 2010 policy change that introduced a tax subsidy limit on annuity pension contributions. Contributions above this new limit could not be deducted in the taxpayer's income. This cutoff is used to split savers in a treatment group (those above the limit) and a control group (those below the limit). Despite their very different identification stra-

tegies and despite the fact that these papers use tax reforms implemented more than a decade apart, they come to the same conclusions; tax incentives for saving in retirement accounts do not increase savings at all. The quasi-experimental nature of these studies and the longitudinal dimension of the data allows the authors to overcome many of the challenges previously discussed in the literature.

The second challenge highlighted by Engen *et al.* (1996) concerns with the fact that household debt is broadly overlooked in crowding-out literature. One of the first papers addressing the possible interplay between household debt and pension contributions is Amromin *et al.* (2007), who showed that some households would be better off by slowing down mortgage debt repayments while raising pension contributions. Here, the authors define an optimal portfolio choice and use US survey data to quantify how many households that fail to optimize the allocation of debt and assets. Despite that this study is not concerned with the effects of tax incentives it is important to mention because it illustrates how complex financial decisions households face; whether to save in retirement accounts, non-retirement accounts or repay debt. Andersen (2018) is the first paper to include both mortgage and non-mortgage debt when assessing the effectiveness of tax incentives for saving in retirement accounts. The next section elaborates on this paper.

3. The pension and debt interplay

The question asked in this section is whether tax incentives for saving in pension accounts cause individuals to cut debt repayments in order to contribute to tax-favored retirement accounts. This would lead to an accumulation of both debt and pension savings, while net savings remain unchanged. In other words, do the tax incentives cause savers to reshuffle their financial portfolios such that they refrain from repaying debt in order to benefit from the tax break on pensions. If this is the case then the beneficial tax treatment of pension contributions could in part explain the increase in debt accumulation over the past decades. Andersen (2018) shows that tax benefits put in place by the government have no effect on overall savings but affect the composition of individual financial portfolios. In fact, the results indicate that such tax incentives cause financial assets and liabilities to move in the same direction, consistent with the so-called balance sheet expansion in households that has been observed in a range of developed countries.

Engen *et al.* (1996) asked for a net savings measure, i.e. total assets minus debt, rather than total assets alone. Their argument is straight forward; households could borrow 1,000 dollars and invest all of them in tax-favored pension accounts, leaving their net savings unchanged while gross savings increase by 1,000 dollars. Their data supported the need for including financial liabilities when analyzing crowding-out effects because households who contributed to retirement

accounts had more debt than the non-contributors. Almost two decades later, Chetty *et al.* (2014) offered the first empirical study that incorporates part of the liability side, namely bank credit. Information about mortgage debt was, however, not available to these authors and given that they were omitting what usually constitutes the largest part of household liabilities, they could not reject that households' adjusted their debt repayments when contributing to pension schemes.

The first empirical study on crowding-out in retirement accounts to incorporate all financial wealth and debt is Andersen (2018). This paper suggests that savers substitute pension savings for debt repayments by 1/3 when tax incentives for saving in pension accounts are reduced. For each 1 dollar reduction in pension contributions debt repayments increased by about 33 cents. The remainder of the dollar is saved in other savings accounts, implying that total savings are unchanged. The Danish tax authorities provide information about individual wealth in bank accounts, stocks and bonds, including a separate measure of outstanding debt to banks. Further, detailed information about outstanding mortgage debt is provided by Finance Denmark, the association of Danish banks and mortgage institutions. It is an important innovation compared to previous papers to be able to follow the cash flow in both assets and liabilities in a population-wide panel data set. The information is audited by government authorities or banks and is not self-reported unlike information from consumer surveys usually used in the literature. This eliminates the risk of self-reporting biases and reduces noise in the data that potentially comes from respondents' inability to remember exactly the size of, e.g., their wealth, debt, income or pension contributions. In addition, the administrative records hold a range of useful information on observables which can be used to filter out effects from e.g. age, gender, educational attainment, employment status and housing.

Access to high quality microdata is not sufficient to examine the effect of pension taxation on savings behavior. A transparent and well designed empirical strategy is equally important. In order to identify the effect of the tax incentives on savings we need savers to be affected differentially in terms of taxation, while everything else remains unchanged. Moreover, we need to ensure that the savers considered in the study had equally high tastes for saving. To overcome these two challenges the paper exploits a Danish 2010 pension tax policy change. This reform introduced a tax deduction threshold of DKK 100,000 (about US \$15,000) on annuity pension contributions. This change affected savers who usually contributed more than DKK 100,000 to this type of pension scheme, while savers who usually contributed less than this cutoff should be unaffected by the policy change. By assuming that tastes for saving and savings preferences in general are close to identical for individuals just above or just below the DKK 100,000 threshold, it is argued that the only difference between the two groups is the tax benefit obtained by saving in retirement accounts. The above-group experienced an unex-

pected reduction in tax deduction when saving in pension accounts, while the below-group did not face any changes in tax treatment. The simple idea of the empirical strategy is to follow and quantify cash flows in all savings and debt accounts undertaken by these two groups and compare how their behavior diverges as the reform is implemented.

The findings show full substitution within savings and debt accounts, implying that overall savings remained unchanged despite the reduction in tax benefits from saving in certain retirement accounts. Moreover, about 1/3 of money usually saved in pension accounts was now used to repay debt, mainly credit to banks. These results are important because they document a causal relationship between pension taxes and accumulation of debt in households. In other words, the rise in both pension assets and liabilities in the household sector could in part be explained by the beneficial tax treatment of savings in pension accounts.

There is, however, more questions to be answered in this literature. In particular three questions come to mind. First, how would savers respond to an increase in pension tax deductions? Andersen (2018) examines a reduction in pension tax benefits and the findings relate only to this type of policy change. It is natural to assume that the response is symmetric such that savers reduce debt repayments when the government increase tax benefits on pensions but this is not answered directly in the existing literature. Another question to ask is whether the causality also exist in the opposite direction, i.e. whether savings in pension accounts would increase if tax deductions on debt interest payments became less favorable. Finally, empirical evidence offered by Chetty *et al.* (2014) and Andersen (2018) has limited external validity in the sense that their conclusions relate mainly to the particular groups of individuals examined in their studies. In other words, findings from quasi-experimental papers should only be extrapolated to the full population with caution when the sample examined covers only a subgroup of the population. For instance, while Chetty *et al.* (2014) examines behavior of individuals from the middle of the income distribution and Andersen (2018) examines behavior of individuals from the top of the income distribution. Empirical evidence on crowding-out effects in retirement accounts remains to be documented specifically for low-income groups.

4. Conclusion

This paper reviewed recent findings in public finance about the interplay between pension savings and debt in households and their link to the macroeconomic developments. The first part deals with the importance of household debt levels on real economic outcomes. Private consumption has remained low for a longer period of time following the 2007–2008 recession compared to earlier economic downturns, potentially because households increased the pace of deleve-

raging which forced them to cut spending. Large debt positions in households are therefore likely to amplify volatility in the economy by causing larger booms and busts. The second part of the paper attempts to explain the high debt levels in household by asking; have increased pension savings caused households to accumulate more debt than they would have done in absence of tax-favored pension accounts?

Findings from Andersen (2018) show a causal link from pension tax benefits on debt accumulation. By exploiting variation caused by a tax policy change that made pension contributions less favorable in terms of taxation the empirical evidence documents an increased pace of deleveraging by the affected households.

Further research is necessary to uncover more details about household savings behavior. Data-demanding tasks such as natural experiments are excellent in terms of quantifying the savings response by households when, e.g., institutional settings or tax rules change. The results from such studies can be used as input in structural models and New Keynesian-type models in order to provide important welfare analysis.

References

- Amromin, Gene, Huang, Jennifer, & Sialm, Clemens. 2007. The Tradeoff between Mortgage Repayments and Tax-deferred Retirement Savings. *Journal of Public Economics*, **91**(10), 2014–2040.
- Andersen, Asger Lau, Duus, Charlotte, & Jensen, Thais Lærkholm. 2016. Household debt and spending during the financial crisis: Evidence from Danish micro data. *European Economic Review*, **89**, 96 – 115.
- Andersen, Henrik Yde. 2018. Do tax incentives for saving in pension accounts cause debt accumulation? Evidence from Danish register data. *European Economic Review*, **106**, 35 –53.
- Attanasio, Orazio P., & DeLeire, Thomas. 2002. The Effect of Individual Retirement Accounts on Household Consumption and National Saving. *The Economic Journal*, **112**(481), 504– 538.
- Attanasio, Orazio P., & Rohwedder, Susann. 2003. Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom. *The American Economic Review*, **93**(5), 1499–1521.
- Benjamin, Daniel J. 2003. Does 401(k) eligibility increase saving? Evidence from propensity score subclassification. *Journal of Public Economics*, **87**(2003), 1259–1290.
- Bernheim, Douglas. 2002. Taxation and Savings. *Chap. 18 of: Auerbach, Alan J., & Feldstein, Martin (eds), Handbook of Public Economics 4*, vol. 3. Elsevier Science.
- Cecchetti, Stephen G., Mohanty, M. S., & Zampoli, Fabrizio. 2011 (September). *The Real Effects of Debt*. Bank of International Settlements Working Papers No. 352.
- Chetty, Raj, Friedman, John N., Leth-Petersen, Søren, Nielsen, Torben Heien, & Olsen, Tore. 2014. Active vs. Passive Decisions and Crowd-out in Retirement Savings Accounts: Evidence from Denmark. *Quarterly Journal of Economics*, **129**(3), 1141–1219.
- Dynan, Karen E. 2009. Changing Household Financial Opportunities and Economic Security. *Journal of Economic Perspectives*, **23**(4), 49–68.
- Dynan, Karen E. 2012. *Is Household Debt Overhang Holding Back Consumption*. Brookings Institute.
- Dynan, Karen E., Skinner, Jonathan, & Zeldes, Stephen P. 2002. The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer. *American Economic Review*, **92**(2), 274–278.
- Eggertsson, Gauti B., & Krugman, Paul. 2012. Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. *Quarterly Journal of Economics*, **127**(3), 1469–1513.

- Engelhardt, Gary V., & Kumar, Anil. 2007. Employer matching and 401(k) saving: Evidence from the health and retirement study. *Journal of Public Economics*, **91**(2007), 1920–1943.
- Engen, Eric M., Gale, William G., & Scholz, John Karl. 1996. The Illusory Effects of Saving Incentives on Saving. *Journal of Economic Perspectives*, **10**(4), 113–138.
- Gale, William G. 1998. The Effects of Pensions on Household Wealth: A Reevaluation of Theory and Evidence. *Journal of Political Economy*, **106**(4), 706–723.
- Gale, William G., & Scholz, John Karl. 1994. IRAs and Household Savings. *American Economic Review*, **84**(5), 1233–1260.
- Hubbard, Glenn R. 1984. Do IRAs and Keoghs Increase Savings. *National Tax Journal*, **1**(37), 43–54.
- Hubbard, Glenn R., & Skinner, Jonathan S. 1996. Assessing the Effectiveness of Saving Incentives. *Journal of Economic Perspectives*, **10**(4), 73–90.
- Jorda, Oscar, Schularick, Moritz, & Taylor, Alan M. 2013. When Credit Bites Back: Leverage, Business Cycles and Crisis. *Journal of Money, Credit and Banking*, **45**(2), 3–28.
- King, M. A., & Dicks-Mireaux, L-D. L. 1982. Asset Holdings and the Life-Cycle. *The Economic Journal*, **92**(366), 247–267.
- Laibson, David. 1997. Golden Eggs and Hyperbolic Discounting*. *The Quarterly Journal of Economics*, **112**(2), 443–478.
- Mian, Atif, Rao, Kamalesh, & Sufi, Amir. 2013. Household Balance Sheets, Consumption and the Economic Slump. *The Quarterly Journal of Economics*, **128**(4), 1687–1726.
- Mian, Atif, Sufi, Amir, & Verner, Emil. 2017. Household Debt and Business Cycles Worldwide*. *The Quarterly Journal of Economics*, **132**(4), 1755–1817.
- Poterba, James M., Venti, Steven F., & Wise, David A. 1995. Do 401(k) contributions crowd out other personal saving. *Journal of Public Economics*, **58**(1), 1–32.
- Poterba, James M., Venti, Steven F., , & Wise, David A. 1996. How Retirement Saving Programs Increase Saving. *Journal of Economic Perspectives*, **10**(4), 91–112.
- Skinner, Jonathan S., & Feenberg, Daniel R. 1990. The Impact of the 1986 Tax Reform Act on Personal Saving. *National Bureau of Economic Research*, Februar.
- Venti, Steven F., & Wise, David A. 1990. Have Iras Increased U. S. Saving?: Evidence From Consumer Expenditure Surveys. *Quarterly Journal of Economics*, **105**(3), 661–698.